Causality and Networks

Mathilde Noual

October 28, 2016

arXiv:1610.08766v1 [cs.OH] 26 Oct 2016

Causality is omnipresent in scientists' verbalisations of their understanding, even though we have no formal consensual scientific definition for it. In Automata Networks, it suffices to say that automata "influence" one another to introduce a notion of causality. One might argue that this merely is an incidental side effect of preferring statements expressed in natural languages to mathematical formulae. The discussion of this paper shows that if this is the case, then it is worth considering the effects of those preferences on the contents of the statements we make and the formulae we derive. And if it is not the case, then causality must be worth some scientific attention *per se*. In any case, the paper illustrates how the innate sense of causality we have may be made deliberate and formal use of without having to pin down the elusive notion of causality to anything fixed and formal that wouldn't do justice to the wide range of ways it is involved in science-making.

1 Boolean Automata Networks

A network is a set of entities/parameters causing each other to undergo change.

A Boolean Automata Network (BAN) $\mathcal{N} = \{f_i : \mathbb{B}^n \to \mathbb{B} = \{0,1\} \mid i \in V = [\![1,n]\!]\}$ is a set of Boolean functions that define the way automata in the set $V = [\![1,n]\!]$ interact with one another.The interaction graph $G = (V, A)$ of $\mathcal{N}$ represents these interactions. Arc $(i,j) \in A$ belongs to the interaction graph G of $\mathcal{N}$ if and only if x_i appears in the Conjunctive Normal Form of $f_j(x)$ – *i.e.* if and only if $f_j(x)$ depends on x_i. When $\mathcal{N}$ is in state $x = (x_1, x_2, \ldots, x_j, \ldots, x_n) \in \mathbb{B}^n$, automaton $j \in V$ has the possibility of undergoing the change of states $x_j \rightsquigarrow \neg x_j$ if it is unstable, meaning: if $j \in U(x) = \{i \in V : f_i(x) \neq x_i\}$. It only actually does undergo the change of states $x_j \rightsquigarrow \neg x_j$ if on top of being unstable in x, j is also *updated* in x (by us).

As mentioned above, in BANs, it suffices to say that automata "influence" one another to introduce a notion of causality. It is very natural to interpret arc $(i,j) \in A$ as meaning "automaton i can cause automaton j to change states". There are other notable examples of causality's spontaneous involvement in the BAN formalism. . .

Additional definitions

Before moving on to the next sections, I will introduce here some additional definitions that the sequel refers to.

The parallel update schedule

The parallel update schedule of a BAN $\mathcal{N} = \{f_i : \mathbb{B}^n \to \mathbb{B} = \{0,1\} \,|\, i \in V = [\![1,n]\!]\}$ is defined by the function $F : (x_1, x_2, \ldots, x_n) \in \mathbb{B}^n \mapsto (f_1(x), f_2(x), \ldots, f_n(x)) \in Bn$. The *transition graph* induced by the parallel update schedule is the graph of function F, namely $T_p = (\mathbb{B}^n, \rightarrow)$ where $(x,y) \in \rightarrow \iff y = F(x)$.

In $x \in Bn$, the parallel update schedule exploits all instabilities of $U(x)$, *i.e.* it exploits all possibility of synchronous updates.

The asynchronous setting

Asynchronism is not an update schedule. It is a notion that can mean $|U(x)| = 1$. It however usually is rather taken to refer to a very popular setting of the literature of (B)ANs. In this setting, the transition graph considered is $T_a = (\mathbb{B}^n, \rightarrow)$, namely the *asynchronous transition graph* where $(x,y) \in \rightarrow \iff \exists i \in U(x) : D(x,y) = \{i \in V : x_i \neq y_i\} = \{i\}$.

In the asynchronous setting, all sequential trajectories are assumed to be possible. All trajectories involving some synchronous updates are assumed not to be.

General transition graphs

The general transition graph of a BAN is the graph $T_g = (\mathbb{B}^n, \rightarrow)$ where the relation $\rightarrow$ is given by: $(x,y) \in \rightarrow \iff \forall i \in V, x_i \neq y_i : i \in U(x)$.

General transition graphs are simply state transition systems. They were introduced in [4] to to study the *behavioural possibilities* of networks as opposed to the *dynamics*.

Block-sequential update schedules

A *block-sequential update schedule* (BSUS) is an update schedule that updates automata of a network in a certain deterministic periodic order and with a certain amount of synchronism in the updates (possibly none). Within a period of updates, all automata are updated exactly once.

Formally it can be defined as a function $\nu : A \to \{-1, +1\}$ such that starting in $x \in \mathbb{B}^n$ at the beginning of the period, $\forall (i,j) \in A$, $\nu(i,j) = -1$ iff i is updated strictly before j is. Thus, when j is updated, i is already in state $f_i(x)$. Otherwise, if $\nu(i,j) = +1$,

then i is updated no sooner than j is so when j is updated, i is still in state x_i.

The "degree of synchronism" of ν might be formally matched to the number of arcs $(i,j) \in A$ such that $\nu(i,j) = +1$. This is only a suggestion however because the definition of BSUS given here is not common.

Traditionally, BSUSs are rather defined as lists of disjoint blocks of automata $(B_k)_{k \leqslant K}$ where $\biguplus_{k \leqslant K} B_k = V$. Automata in a block are updated in parallel. The blocks are updated sequentially[5,6]. With this definition, a BSUS has more synchronism if it has less and larger blocks. For example consider the BAN of Fig.1. It is what we call a Boolean Automata Cycle (BAC). One particular BSUS of this BAC is the following which defines $K = n$ blocks of size 1: $\{n\}, \{n-1\}, \ldots, \{3\}, \{2\}, \{1\}$. This BSUS sequentially updates automata in the reverse order of the cycle. It has no synchronism. Another BSUS of the BAC is: $\{6, \ldots, n-1, n\}, \{2, 3, \ldots, 5\}, \{1\}$ which defines $K = 3$ blocks. Two of them are of size greater than 1 so this BSUS *does* have some synchronism. It updates at once all automata in block $B_1 = \{6, \ldots, n-1, n\}$. Then it updates at once all automata in block $B_2 = \{2, 3, \ldots, 5\}$. And only then does it update automaton 1. The BSUS $\{2, 3, \ldots, n-1, n\}, \{1\}$ has even more synchronism since it updates at once all automata except automaton 1, before it updates automaton 1.

Notably, all three of these BSUSs are equivalent to $\nu : A \to \{-1, +1\}$ defined by $\nu(i,j) = 1 \Leftrightarrow (i,j) \neq (n,1)$.

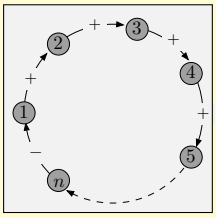

$$\begin{cases} f_i : x \mapsto x_{i-1} & \forall i \neq 0 \\ f_1 : x \mapsto \neg x_n \end{cases}$$

Fig. 1: A negative *Boolean Automata Cycle*, a.k.a. negative *feedback loop*. LEFT Its interaction graph. RIGHT: Its defining local update functions. The BAC owes its negativity to the odd number of negative arcs in the interaction graph.

2 Synchronism *vs* Precedence

In some circumstances, a synchronous updating of automata seems to cause local instabilities to linger in the network longer than they might with more sequentiality. In other terms, synchronism seems to inhibit the decrease of the number $|U(x)|$ of automata that are unstable. And as a consequence, it also seems to be responsible for the global asymptotic instability of the BAN [1,2–4].

Indeed, on the one hand, the BSUSs of BACs that allow for the most local and global instability are the BSUSs with a less synchronism [1,2]. Moreover, we know that with the parallel update schedule a BAC has many attractors. And the majority of those attractors involve several local instabilities circulating through the network [7] . On the

other hand, in the asynchronous setting, all BAC trajectories lead to an attractor that has either 0 or 1 instabilities[8]. And the general transition graphs of BACs show that at least some sequentiality is required at some point to reduce the number of instabilities in the BAC.

All this supports the idea that synchronism is responsible for entertaining instabilities. In communities focusing on asynchronous Automata Networks to model genetic regulation networks, synchronism is actually notorious for this effect. Synchronism's alleged tendency to artificially entertain instabilities in a network is a very common informal argument invoked to undermine the realism of updating schedules involving synchronism, and the relevance of research that doesn't rule synchronism out altogether. As I was cutting my teeth on such research [2,3,7,9–14] , it was necessary for me to investigate the causal link between synchronism and instability [15] .

As it turns out, when a small amount of synchronism is added to an otherwise asynchronous BAN, synchronism either has no lasting effect (most frequent case by far), or it has some (this is the case with the BAN of Fig.2). And if it has some then, precisely, its effect is to stabilise all local instabilities. And in that, it stabilises the whole BAN [15] (see Fig.3).

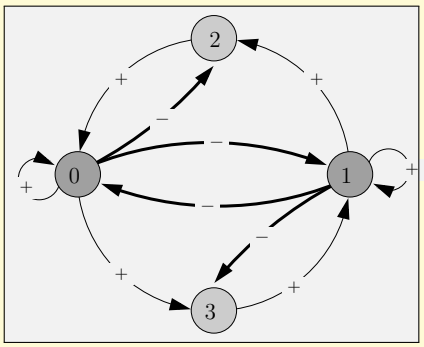

$$\begin{cases} f_0 : x \mapsto x_2 \vee (x_0 \wedge \neg x_1) \\ f_1 : x \mapsto x_3 \vee (\neg x_0 \wedge x_1) \\ f_2 : x \mapsto \neg x_0 \wedge x_1 \\ f_3 : x \mapsto x_0 \wedge \neg x_1 \end{cases}$$

Fig. 2: This is the BAN given in [15] as an example of "synchronism-sensitive" BAN. Its asynchronous transition graph T_a is represented and discussed in Fig.3 together with the addition of some synchronism to T_a.

So on the one hand, synchronism seems to cause instabilities to linger, on the other it 'causes' instabilities to disappear. The quotes around 'causes' are to help keep in mind that the notion of causality we are manipulating here is not one that is formally defined. In both situations, synchronism is seen as a cause of something, but in none is it actually the logical implicant of anything. Causality, which is not logical implication, is much more charged with meaning than logical implication is. Here, it is charged in particular with the specific meaning of the term "synchronism" that is used in each of the two situations (*cf.* paragraph on Block sequential update schedules on Page 3).

The conclusion we can draw from the apparent contradiction is the following. We have taken a certain perspective. According to this perspective, 'synchronism' is a

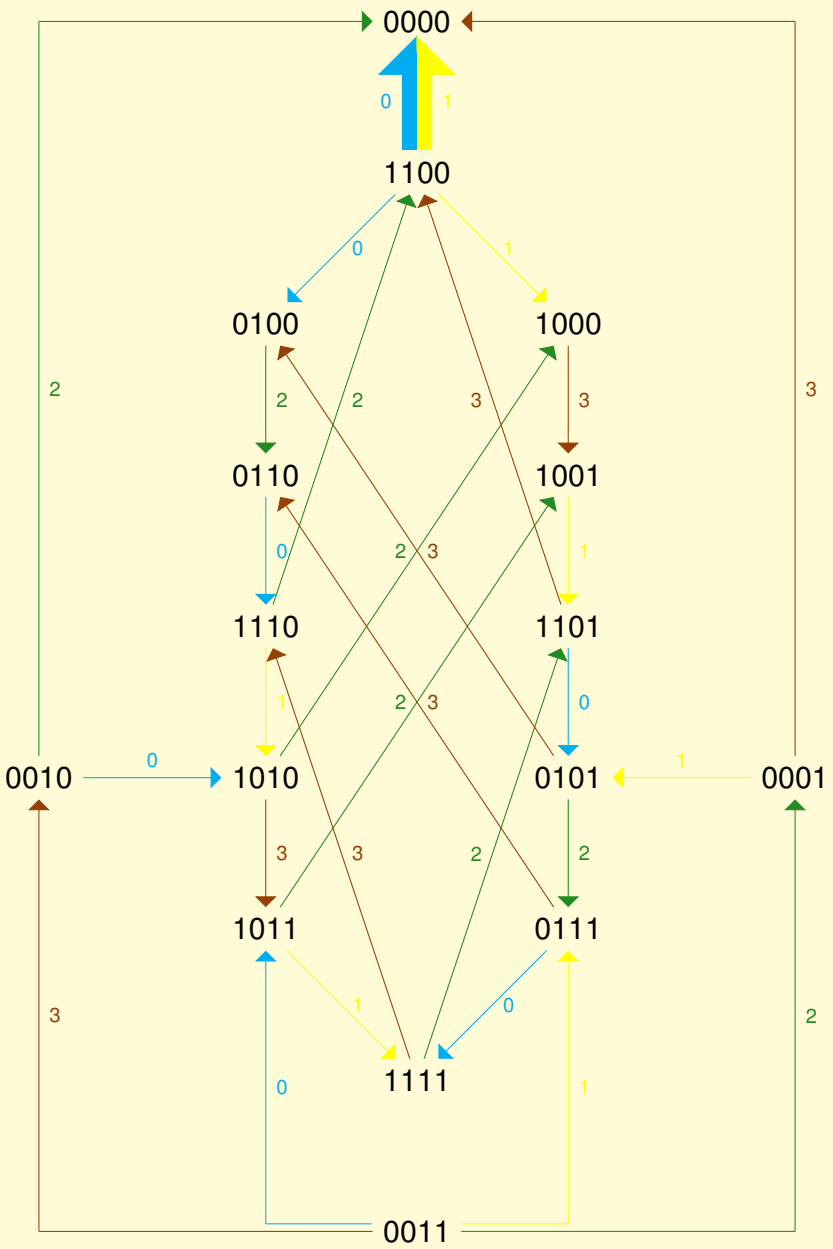

Fig. 3: The asynchronous transition graph of the BAN of Fig.2 with the additional synchronous transition $1100 \xrightarrow{01} 0000$. All other synchronous transitions of the BAN do not have a lasting effect: they merely shortcut asynchronous trajectories. With a purely asynchronous updating of automata states, the BAN has two attractors: (1) the stable state 0000 ($U(0000) = V$) and (2) a large cyclic attractor comprised of the 12 states at the centre. When the transition $1100 \xrightarrow{01} 0000$ is added, the BAN looses its second attractor. The thing to notice is the following. Consider the state $x = 1100$. In this state, if the two instabilities – that of automaton $0 \in U(x)$ and that of automaton $1 \in U(x)$ – are not settled – *i.e.* if the synchronous transition is not made – then either the two parts of what constitutes a XOR $x_0 \oplus x_1$, namely $(x_0 \wedge \neg x_1)$ and $(\neg x_0 \wedge x_1)$ start being schlepped around in the network, from one automaton to the next, thereby entertaining a certain tacit dependency between the instabilities of automata, precisely the dependencies that characterises the second attractor of the asynchronous updating. It isn't yet clear however whether or not they assemble straightforwardly into $x_0 \oplus x_1$ at some point. But the BAN of Fig.2 being a minimal example of synchronism-sensitive BAN [15] , this suggests there is a relation between synchronism-sensitivity and non-monotony.

meaningful notion: it is a notion that pinpoints actual causes of notable effects we are interested in. More specifically, according to this perspective, 'synchronism' is the type of thing that can cause an effect on local instabilities. With this perspective, we run into an apparent contradiction. *Two opposite effects are incumbent on the same cause.* And that in itself is a call for an update of perspective. Synchronism (and asynchronism for that matter) might not be the appropriate notion to explain either of the effects we want it to explain in relation to instabilities.

Nonetheless, the intuitions expressed by the two different cause/effect relations involving synchronism and instabilities cannot be baseless. Or at least this much is true: considering that they are baseless is not going to help us make any progress, we had better trust that they aren't and go looking for their basis in order to formalise it. Ruling out synchronism as a possible medium for the effects of entertaining and settling instabilities means that we need a new candidate cause that can convey a finer form of causality accounting coherently for both effects. Comparing the two contradictory situations reveals that *precedence* might be a more reasonable choice for that. Just like synchronism, it is involved in both situations. And it is a plausible basis for both sides of the contradiction as it can explain both situations coherently. Indeed, in the first situation, instead of "*Synchronism causes instabilities to linger*" or rather instead of "*In the absence of synchronism, instabilities do not linger (as much)*", it seems more relevant to use the following explanation of what happens to instabilities:

"*Assuming event B owes its possibility to event A not having occurred, then, in case of precedence of the occurrence of possible event A over the occurrence of possible event B, the occurrence of A causes the impossibility of the occurrence of B.*"

meaning that the implementation of instability A (*i.e.* the possibility of A) is enough to settle at once both instability A and instability B. And in the second situation, instead of "*Synchronism settles instabilities*", we can use the following explanation:

> "*If two events A and B are possible, if the occurrence of each results in the disappearance of one instability, then in the absence of any interference of precedence, the occurrence of A and B results in the disappearance of two instabilities.*"

The effects studied in [15] were classified in terms of varying degrees of *sensitivity to synchronism*. But following this discussion, *sensitivity to precedence of causally related events* might be a more accurate and relevant way of coining the same effects.

Generally, if $(i, j) \in A$ and $i \in U(x)$, $j \notin U(x)$, then updating i before j – assuming j owes its stability in x to i being in state x_i – causes i to stabilise and j to become unstable in turn. If both automata $i \in U(x)$, $j \in U(x)$ are unstable in x, then – assuming

x_i is the reason for j's instability in x – updating i before j will stabilise i in state $\neg x_i$ and thereby also cause the stabilisation of j in state $x_j = f_j(\bar{x}^i)$.

It no longer is a matter of synchronism.

On top of putting forward precedence as a more relevant and accurate cause, the comparison of the two initial contradictory cause/effect relations emphasises all the other differences there are between the two situations.

Determinism of the parallel update schedule and of BSUSs is one of them. And this difference with the asynchronous setting relates to the notable fact that asynchronism is not an update schedule, and it is thereby not comparable with sequential update schedules. T_a, just like T_g, defines a state transition system representing *behavioural possibilities*, as opposed to *defined dynamical behaviours*. It is not clear how to compare what is called an "attractor" in the asynchronous setting (the terminal strongly connected components of T_a) to the attractors of dynamics defined by specific deterministic update schedules.

Periodicity – perhaps even more subtly: the specific kind of redundancy inherent to periodicity – is another notable difference between the two situations. In the lead of F. Robert[5], a great many studies have been supporting the general idea that "update schedules have great influence on the dynamics of BANs" [16–18, 2, 3, 14, 19]. But just like precedence and determinism, until periodicity's own effects aren't studied *per se*, there is no rigorous way to form a more reliable intuitive understanding of what generally causes the entertainment of local instabilities under BSUSs and of what, other than asynchronism, tends to prevent the entertainment of local instabilities when synchronism is not exploited.

3 Interlude

The previous section discussed a case where two different effects are intuitively attributed to the same cause in different contexts, and shows how to learn from that by letting the following question be raised: *How is the cause really involved in the generation of the effect?* and letting it lead to the next question, namely *What else is involved? in particular, what finer cause might the original cause be a facade for or an abstraction of?*

Another case we can learn from is when two different causes are intuitively invoked to explain the same effect in different contexts. In this case the following question calls for an answer: *What is the effect's common implicant/generating mechanism?* This case is illustrated in the next section.

4 Oscillations and the difference between experience of change and communication of change

In some contexts, negative feedback loops, or negative cycles (*cf.* Fig.1) are considered to be directly responsible for the asymptotic oscillations of BANs [20–24]. In different contexts, BANs can have asymptotic oscillations without negative cycles [16, 25, 7, 26]. Asymptotic oscillations can therefore not intuitively be attributed to negative cycles (nor to the specific context of each case for that matter). Negative cycles are not the *fundamental* generating mechanism of asymptotic oscillations. A close comparison between those apparently contradictory cases reveals however, that with or without negative cycles, asymptotic oscillations are essentially generated the same way: "something" is disallowing the collapse of an offset between the actual state of a certain automaton i and the pending influence it sends out to itself possibly via other automata. In the first context, that "something" is a negative feedback loop. In the other contexts, it is a combination of a positive feedback loop, some *in situ* potentiality (defined by $x \in \mathbb{B}^n$), and some ~~synchronism~~ absence of precedence. In those other contexts the same effect is produced with what can be argued to be less numerous and less elaborate means, requiring neither precedence nor negation, just the feedback loop and *in situ* information.

Now, we have identified a common generating mechanism with two possible implementations explaining the same effect in two different contexts. It remains to address the questions that naturally follow starting with *What other ways are there to implement this mechanism?* and *Can this mechanism be held responsible for other effects?*

5 Interlude

The way Sections 2 and 4 suggest to make use of the notion of causality is by considering it as essentially progressive so that we let causal relations point towards the signs that already exist of their own coming obsolescence. A particular advantage of taking such a flexible, yet deliberate and rigorous perspective on causality is that it keeps us aware of the possibility of having different causes causing the same effects and thus that the explanations we currently have of the effects we currently are interested in might still be worth investigating even if they appear to be perfectly functional and complete explanations. This perspective on causality also prepares us to deal with new situations in which we observe the same familiar effects in a system, without having any formal reason to believe that the same implicants we are used to are responsible. And it favours keeping in mind the possibility that causes we have represented of the effects we have observed, can themselves be abstractions of more subtle, atomic mechanisms operating at a lower level of abstraction. It encourages finding ways to exploit this possibility in order to refine our explanations and understanding. Generally, this approach to causality emphasises the fact that for a

significant part, what we manipulate as scientists is representations of the objects that we study, not just the objects themselves. And in emphasising this fact, it allows us to take advantage of it.

Now of course, BANs can be studied with purely mathematical interests and perspectives. Then, causality is not such an important concern. Implication is enough. But because it applies only to specific properties of specific systems in specific conditions, implication has the downside of being much less portable than causality. Moreover, in theory, we are free to pick any restriction on the kinds of BANs we consider. And often, when BANs are studied for purely mathematical reasons, the priority is to pick a restriction that will help derive new mathematical results ...

6 Non-monotony and efficiency

A common restriction motivated by biological considerations is the restriction to monotone (B)ANs [8,21,27–29]. Fig.4 shows that a monotone BAN can behave just like a non-monotone BAN in some conditions (see also [30]). Conversely, Fig.5 shows that a non-monotone BAN can behave just like a monotone BAN in some conditions. Studying monotone BANs and non-monotone BANs is therefore not enough. Even comparing them is not enough because at most what we get from doing that is a list of monotone BAN properties that non-monotone BANs don't share, and a list of non-monotone BAN properties that monotone BANs don't have. If a monotone BAN can behave exactly like non-monotone BAN and *vice-versa*, then what we need to know, is how *non-monotony* and *monotony* work: *What are the effects/properties each of them is strictly responsible for?*

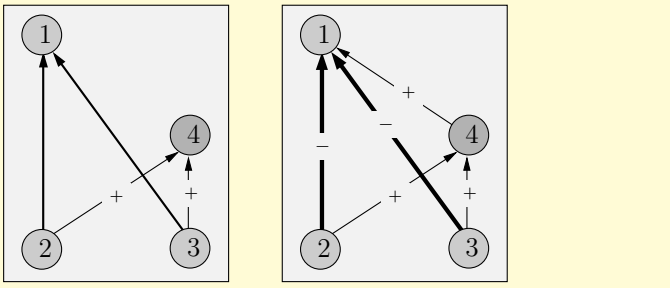

$$\begin{cases} f_1 : x \mapsto x_2 \oplus x_3 \\ f_4 : x \mapsto x_2 \vee x_3 \end{cases} \qquad \begin{cases} f_1 : x \mapsto (\neg x_2 \vee \neg x_3) \wedge x_4 \\ f_4 : x \mapsto x_2 \vee x_3 \end{cases}$$

Fig. 4: LEFT: Non-monotone interactions with automaton 1. RIGHT: Monotone interactions with automaton 1 that can produce exactly the same effects if the right relative arrangement in time of automata updates is chosen.

As illustrated in Fig.4, for a monotone BAN's behaviour to look just like that of a non-monotone BAN, it requires the right arrangement in time of automata updates, and it requires us ignoring certain *intermediary steps*. For the monotone BAN on the right of Fig.4 to behave like the non-monotone BAN on the left, automaton 4 needs to be

updated systematically before automaton 1 is, and the two sequential transitions $x = (x_1, x_2, x_3, x_4) \xrightarrow{4} x' = (x_1, x_2, x_3, f_4(x)) \xrightarrow{1} x'' = (f_1(x'), x_2, x_3, f_4(x))$ need to be summed up into just one transition $x = (x_1, x_2, x_3, x_4) \longrightarrow x'' = (f_1(x'), x_2, x_3, f_4(x))$. Incidentally, since automaton 4 no longer has any impact, we can loose it altogether and concentrate on the asynchronous transition $(x_1, x_2, x_3) \longrightarrow (f_1(x'), x_2, x_3)$. In any case, the example of Fig.4 shows that the intrinsic effects of non-monotony that cannot be reproduced without non-monotone f_i's, *i.e.* effects that represent the real difference between the two types of BANs, are strongly related to time flow (the relative order of events considered) and to the degree of precision of our observations of global trajectories (with more or less intermediaries). This suggests that non-monotony *per se* might impact in terms of the efficiency of the execution of mechanisms in BANs.

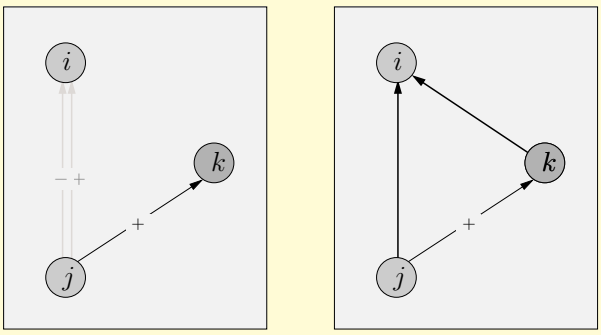

$\begin{cases} f_i : x \mapsto x_j \oplus x_j = 0 \\ f_k : x \mapsto x_j \end{cases}$ $\begin{cases} f_i : x \mapsto x_j \oplus x_k \\ f_k : x \mapsto x_j \end{cases}$

Fig. 5: LEFT and RIGHT, the function defining how automaton i behaves is $f_i(x) = x_j \oplus x_\ell$ for $\ell \in \{j, k\}$. Whether or not j is considered to be an influencer of i depends on whether $\ell = j$ or $\ell \neq k$. And yet, the non-monotone BAN on the right can be made to behave exactly as the monotone BAN on the left by imposing the systematic precedence of the update of k over that of i so that $\forall x,\ f_i(x) = x_j \oplus f_k(x) = xj \oplus x_j$.

Of course, such a claim needs to find some formal support. Even though the notion of causality it conveys is unlikely to be exhausted in one proven mathematical statement, the claim might nonetheless be supported to some extent by putting forward proven mathematical statements that are compatible with it just like the example of Fig.4 is. But Fig.5 suggests that just like with synchronism, at some point, we might need to call into question non-monotony's capacity at explaining relevantly the effects we attribute the responsibility of to it. And together with that, we'll have to call into question the perspective that places the focus on the distinction monotony/non-monotony rather than on a more subtle distinction, one that can provide more complete explanations with less means. In other terms, non-monotony will eventually have to give way to a finer explanation of the same effects. And considering Fig.5, it seems that two notions implied by non-monotony might then be put forward instead of non-monotony: a notion of "witness" and a notion of "inconsistent parallel transfer of the same information".

7 Non-expansivity

In some (asynchronous) contexts[20,31], a BAN $\mathcal{N}$ is defined using the global update function $F : x \mapsto (f_1(x), \ldots, f_n(x)) = \mathcal{N}$ which happens to coincide with the function defining the dynamics of $\mathcal{N}$ under a parallel update schedule. The definition is equivalent to the one chosen in Section 1, namely $\mathcal{N} = \{f_1, \ldots, f_n\}$. But it makes it more natural to pick restrictions on BANs that are given by properties of F. An example is non-expansivity: $\forall x, y \in \mathbb{B}^n : |\{x_i \neq y_i\}| \geqslant |\{F(x)_i \neq F(y)_i\}|$. Intuitively, F's non-expansivity corresponds to the BAN having a form of global instantaneous potential. Assuming F's non-expansivity happens to favour the derivation of some results about asynchronous BANs [32,33]. But precisely, asynchronous BANs are BANs whose dynamical constraints forbid them the use of this global potential. This makes it difficult to grasp intuitively the meaning of the results in question. And mathematical results that are more difficult to have an intuitive grasp on are often results that are more difficult to generalise and relate to other results.

8 Assumptions and intuitions

The examples of sections 6 and 7 show that however mathematically sound are the mathematical results we prove thanks to mathematical assumptions/restrictions, disregard for intuitive causality at worst stakes the applicability of those results, and generally limits our progression. In particular, without deliberate care, there is no reason to believe that we owe the deriving of these results to some deeper opportune relevance of the mathematical assumptions/restrictions. There is no reason to believe there is anything in the assumptions/restrictions that could enable the generalisation of the results beyond the setting they define, nor anything that could at least guarantee the relatability of the results to other existing results. The primary reason why we might have managed to derive anything under a particular assumption/restriction might be that it is an extremely strong assumption/restriction. It might be like studying crows by concentrating on the class of crows that a human being has reported seeing picking up a piece of pink plastic wrapper. It might be quite unclear what it is that we are studying and learning about exactly: the original (mathematical) object of interest? the restriction? And in the case of Automata Networks, this means that the only hope to actually build a global understanding of networks lies in the platonic wager that it will necessarily "emerge" from the accumulation of independent studies made of particular models of networks in different settings, sometimes juxtaposed for comparison. In lines with Section 2, let us mention as an example that a great amount of rigorous theoretical attention has been invested in the study of a large variety of update schedules [5,6,16–18,34–39,2,3,14,19]. And yet, none of the studies it occasioned, not even the ones that *compare* specific update schedules, have yet provided explicit answers to the kind of fundamental questions mentioned in Section 3, including questions about

the effect *per se* of the redundancy inherent to the periodicity of a periodic update schedule, the effect *per se* of the determinism of a deterministic update schedule, the effect *per se* of varying amounts of instabilities, and of varying degrees of imposed precedence between their exploitation. This can however naturally be made up for by sometimes (a) focusing primarily on general – *i.e.* fundamental and portable – network attributes (*e.g.* synchronism, non-monotony, reversibility, subsequence) rather than on network-specific properties (*e.g.* particular interaction digraphs, particular f_i's) or restrictions, and by (b) studying rigorously the involvement of those attributes in the behavioural possibilities of networks (as opposed to describing formally a specific dynamics). This way, what we know of networks and what we don't can be re-examined and clarified.

9 Time flow and potential storage/computation space

The systems we consider as scientists are often assumed to be conditional to properties of time (flow). Despite this, their formal definitions sometimes allows properties of time flow to mingle or overlap with their own properties (*cf.* remark made in [30], Section 2, about update constraints and the dynamical systems view on BANs). Thus, time flow contributes to the way we attribute properties to a system. Yet, it isn't always clearly distinguished from causality.

Generally, when it isn't regarded as a pre-existing constraint on the systems we consider, time flow – or rather just "Time" – is seen as a "resource", suggesting that we have a tacit obligation to use it sparingly, and, without fail, in finite quantity. Whatever we call it, we tend to assume that it pre-exists both the systems we study and the attention we invest in them, and that it frames both the systems' behavioural possibilities and the leeway we take on them. Operational Research works its satisfaction and optimisation problems around it; Bio-informatics builds models out of what it knows of it or despite what it doesn't [40–44]; and even for Concurrency, time flow is mostly something *in* which distributed pieces of computation might be reunited [45,46]. Despite its very dense presence in the scientific landscape, time flow is seldom a *primary* object of our attention. This results in plenty space for spontaneous interpretations to operate. And in particular it gives carte blanche to a natural tendency we have to expediently distinguish, confuse, overlook or classify issues and properties related to time flow such as simultaneity, synchronicity, precedence, subsequence, difference, determinism, periodicity, causality, time scale, change, process of change, realism, duration. As a result, we miss out on (a) the vicariance of some of these properties for which time flow might actually not be the exclusive medium, and (b) possible leeway through this vicariance.

In an isolated network, time flow is precisely determined by the set of all possible events considered in the network and by their relative order[1].

Giving to causality the attention advocated for in the previous sections leads to putting emphasis on the notion of time flow. It requires to systematically isolate the involvement time flow has in the effects we take interest in, and decompose it into more atomic properties that can more tightly be held responsible for the effects. Thus, the effects of specific properties of time flow (possibly implicitly assumed by the specific framework, *e.g.* precedence of certain event occurrences over certain others) may be compared with effects of properties specific to networks (properties of the f_is). And thus, we can work towards a better understanding of network sensitivity to time flow. We can start clarifying the kind of information time flow is a vehicle of, as well as the part of the information encoded in the clockworks of networks (the f_i's) that can equivalently be encoded into time flow.

Eventually we may also clarify the kind of computation complexity time flow can manage. Indeed, results [15] suggest that under very specific conditions, attributes of time flow might participate in the overall network computation in ways that are comparable to logical gates.

As argued in [30], the notion of synchronism in BANs is often confused with a notion of simultaneity in 'reality' referring to dates represented by the same point on The Time-Line [30] . Nonetheless, in BANs, the possibility of synchronism is an atemporal relation between possible events. It relates events of the set $\{x_i \rightsquigarrow \neg x_i \,|\, i \in U(x)\}$. Events that are synchronously possible are independently caused: they can't be the cause of one another since they haven't yet occurred. So synchronism implies nothing about time flow. Precisely, the occurrences of the events it relates are not ordered relatively to one another – or at least, in the case of asynchronous BANs, not *yet* in x. If BANs are supposed to model what we observe of real systems, then the implications of this are noteworthy. *How the occurrences of synchronously possible events actually end up arranging themselves relatively to one another* is information. And it is information that is *not* modelled. We can either claim that we have the information (*cf.* assumptions mentioned on page 9 of [30]), or we can look for it. If we claim we have the information, then as [30] demonstrates, we still cannot model it with BANs: we can only *interpret* a feature of BANs as representing it which is very different. To explain a particular ordering of occurrences of synchronously possible events requires certain care in how the notion of causality is handled. The model (*i.e.* the BAN) accounts for the synchronous possibility of the events, and not the relative arrangement of their occurrences. If the model represents what we know of the real system at a certain

[1] Note how different this is from saying that in a network, all possible events considered happen *in* Time and their order is defined *by* Time or relatively to it. In particular, it informally implies a different sense of the term "isolated".

level of abstraction, then the information we are looking for – about the relative ordering of synchronously possible events – cannot be found at the same level of abstraction, or else, we did a bad job in modelling the real system and/or in choosing a formalism to do that. Building a model with the capacity of revealing causal relationships that are otherwise not apparent to us, is not the same thing as building a representation of the set of effects we are interested in. The relative ordering of the occurrences of synchronously possible events creates new relationships between those events – relationships that otherwise don't exist as evidenced by the synchronous possibility of these events. Thus, it brings together otherwise presently independent pieces of information and can cause them to interact as they wouldn't have if they had been left independent. In the case of the BAN of Fig.2, the effect of adding this relation evokes that of logical connectors ($\wedge, \vee, \oplus$, *cf.* caption of Fig.3). Of course, again, this claim begs to be investigated.

10 Conclusion

There is not much reason to believe that there is a scientific definition of the notion of causality that can be found to account completely for the way causality serves science-making. Centuries of modern science-making haven't been enough to find anything close to a satisfying proposition. The notion is much too large and diverse to be fitted exhaustively into a fixed predefined formalisation of it. Besides, causality stands for humans' instinctive way of grasping the world. It is an essential part of our motivation to explore the world further and have more of it grasped. It makes sense to have it serve science-making in an unformalised intuitive way. This is however not a reason for letting it serve science-making in a fortuitous way. And we have shown that rather than sidelining the manifestations of this intuitive instinct of ours, we can advantageously supervise its interference with the scientific formalism we use, and let it serve as a pointer towards knowledge in need of further explicit formalisation.

Bibliography

[1] A. Elena. *Robustesse des réseaux d'automates booléens à seuil aux modes d'itération. Application à la modélisation des réseaux de régulation génétique*. PhD thesis, Université Joseph Fourier – Grenoble, 2009.

[2] E. Goles and M. Noual. Block-sequential update schedules and Boolean automata circuits. In *AUTOMATA*, pages 41–50. Discrete Mathematics & Theoretical Computer Science (DMTCS), 2010.

[3] E. Goles and M. Noual. Disjunctive networks and update schedules. *Advances in Applied Mathematics*, 48(5):646 – 662, 2012.

[4] M. Noual. General transition graphs and Boolean circuits. Research report, 2010.

[5] F. Robert. *Discrete iterations: a metric study*. Springer Verlag, 1986.

[6] J. Aracena, E. Goles, A. Moreira, L. Salinas. On the robustness of update schedules in Boolean networks. *Biosystems*, 97:1–8, 2009.

[7] J. Demongeot, M. Noual, and S. Sené. Combinatorics of Boolean automata circuits dynamics. *Discrete Applied Mathematics*, 160:398–415, 2010.

[8] É. Remy, B. Mossé, C Chaouiya, and D. Thieffry. A description of dynamical graphs associated to elementary regulatory circuits. *Bioinformatics*, 19:ii172–ii178, 2003.

[9] M. Noual. On update schedules and dynamics of Boolean networks. Master's thesis, M1, École normale supérieure de Lyon, 2008.

[10] M. Noual. On the dynamics of two particular classes of Boolean automata networks: Boolean automata circuits and OR networks. Master's thesis, M2, École normale supérieure de Lyon, 2009.

[11] J. Demongeot, H. Ben Amor, A. Elena, P. Gillois, M. Noual, and S. Sené. Robustness in regulatory interaction networks. A generic approach with applications at different levels: physiologic, metabolic and genetic. *International Journal of Molecular Sciences*, 10:4437–4473, 2009.

[12] J. Demongeot, E. Goles, M. Morvan, M. Noual, and S. Sené. Attraction basins as gauges of robustness against boundary conditions in biological complex systems. *PLoS One*, 5:e11793, 2010.

[13] J. Demongeot, A. Elena, M. Noual, S. Sené, and F. Thuderoz. "Immunetworks", intersecting circuits and dynamics. *Journal of Theoretical Biology*, 280:19–33, 2011.

[14] M. Noual. Updating Automata Networks. PhD thesis, École normale supérieure de Lyon, 2012.

[15] M. Noual. Synchronism *vs* asynchronism in Boolean networks. Research report, 2011.

[16] A. Elena, J. Demongeot. Interaction motifs in regulatory networks and structural robustness. In *CISIS*, pages 682–686, 2008.

[17] E. Goles, L. Salinas. Comparison between parallel and serial dynamics of Boolean networks. *Theor. Comp. Sci.*, 396:247–253, 2008.

[18] P.T. Tosic. Cellular automata communication models: Comparative analysis of parallel, sequential and asynchronous CA with simple threshold update rules. *IJNCR*, 1(3):66–84, 2010.

[19] J. Aracena, É. Fanchon, M. Montalva, and M. Noual. Combinatorics on update digraphs in Boolean networks. *Discrete Applied Mathematics*, 159:401–409, 2011.

[20] É. Remy, P. Ruet, and D. Thieffry. Graphic requirement for multistability and attractive cycles in a Boolean dynamical framework. *Advances in Applied Mathematics*, 41:335–350, 2008.

[21] F. Blanchini, E. Franco, and G. Giordano. A structural classification of candidate oscillators and multistationary systems. *bioRxiv*, 2013.

[22] R. Thomas. On the relation between the logical structure of systems and their ability to generate multiple steady states or sustained oscillations. *Springer Series in Synergetics*, 9:180–193, 1981.

[23] E. H. Snoussi. Necessary conditions for multistationarity and stable periodicity. *J Biol Syst*, 6:3–9, 1998.

[24] H. Siebert. Analysis of Discrete Bioregulatory Networks Using Symbolic Steady States. *Bull. Math. Biol.*, 73(4):873–898, 2010.

[25] J. Demongeot, J. Aracena, F. Thuderoz, T.-P. Baum, and O. Cohen. Genetic regulation networks: circuits, regulons and attractors. *Comptes Rendus Biologies*, 326:171–188, 2003.

[26] M. Noual. Dynamics of circuits and intersecting circuits. In *Language and Automata Theory and Applications (LATA)*, volume 7183 of *Lecture Notes in Computer Science (LNCS)*, pages 433–444. Springer-Verlag, 2012.

[27] C. Chaouiya, E. Remy, P. Ruet, and D. Thieffry. *ICATPN*, chapter Qualitative Modelling of Genetic Networks: From Logical Regulatory Graphs to Standard Petri Nets, pages 137–156. Springer Berlin Heidelberg, 2004.

[28] J. Aracena, M. González, A. Zuniga, M. A. Mendez, and V. Cambiazo. Regulatory network for cell shape changes during Drosophila ventral furrow formation. *Journal of Theoretical Biology*, 239(1):49 – 62, 2006.

[29] L. Mendoza and E. R. Alvarez-Buylla. Dynamics of the genetic regulatory network for Arabidopsis thaliana flower morphogenesis. *Journal of Theoretical Biology*, 193:307–319, 1998.

[30] M. Noual. Perspectives on Interaction Systems, and Boolean Automata Networks. Research report, 2016.

[31] J. Aracena, A. Richard, and L. Salinas. Number of fixed points and disjoint cycles in monotone Boolean networks. `http://arxiv.org/abs/1602.03109`, 2016.

[32] A. Richard. A fixed point theorem for Boolean networks expressed in terms of forbidden subnetworks. In *AUTOMATA*, 2011.

[33] J.-P. Comet, M. Noual, A. Richard, J. Aracena, L. Calzone, D. Demongeot, M. Kaufman, A. Naldi, E.H. Snoussi, and D. Thieffry. On circuit functionality in Boolean networks. *Bulletin of Mathematical Biology*, 75(6):906–19, 2013.

[34] E. H. Snoussi. Qualitative dynamics of piecewise-linear differential equations: a discrete mapping approach. *Dynamics and Stability of Systems*, 4:565–583, 1989.

[35] A. Garg, A. Di Cara, I. Xenarios, L. Mendoza, and G. De Micheli. Synchronous Versus Asynchronous Modeling of Gene Regulatory Networks. *Bioinformatics*, 24(17):1917–1925, 2008.

[36] H.S. Mortveit, C.M. Reidys. Discrete, sequential dynamical systems. *Discrete Math.*, 226:281–295, 2001.

[37] C. M. Reidys. Sequential dynamical systems over words. *Annals of Combinatorics*, 10:481–498, 2006.

[38] E. Goles and S. Martínez. *Neural and automata networks: dynamical behaviour and applications*. Kluwer Academic Publishers, 1990.

[39] L. Tournier. *Étude et modélisation mathématique de réseaux de régulation génétique et métabolique*. PhD thesis, Université Joseph Fourier - Grenoble, 2005.

[40] G. Batt, R.B. Salah, and O. Maler. On timed models of gene networks. *FORMATS*, 4763:38–52, 2007.

[41] H. Klarner, H. Siebert, A. Bockmayr. Parameter Inference for Asynchronous Logical Networks Using Discrete Time Series. In *CMSB*, pages 121–130. ACM, 2011.

[42] J.-P. Comet, J. Fromentin, G. Bernot, and O. Roux. A Formal Model for Gene Regulatory Networks with Time Delays. In *CSBio*, pages 1–13. Springer Berlin Heidelberg, 2010.

[43] H. Siebert and A. Bockmayr. Temporal constraints in the logical analysis of regulatory networks. *Theoretical Computer Science*, 391:258–275, 2006.

[44] R. Thomas and M. Kaufman. Multistationarity, the basis of cell differentiation and memory. ii. logical analysis of regulatory networks in terms of feedback circuits. *Chaos*, 11:180–195, 2001.

[45] V. Diekert and G. Rozenberg. *The Book of traces*. Singapore River Edge, N.J. World Scientific, 1995.

[46] Z. Manna , A. Pnueli. *The temporal logic of reactive and concurrent systems : specification*. Springer, 1992.